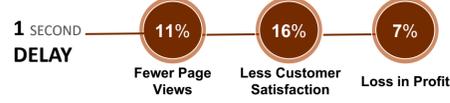


Hyperloop: Improving the Performance of Database-Backed Web Applications

Performance of web applications is important.

- Most users expect a page to load within 2 seconds

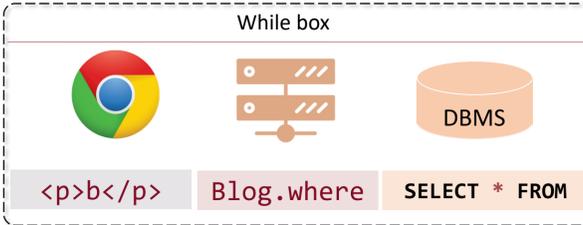


Performance of real-world apps is poor.^[1]

- Popular web apps all have pages taking more than 2 seconds to load
- Server performance is the bottleneck
- Things are getting worse with more data to manage and process

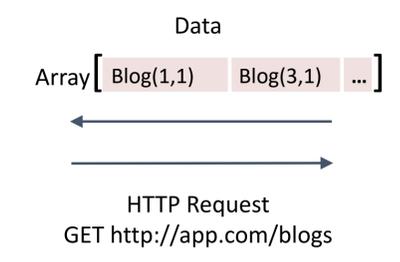
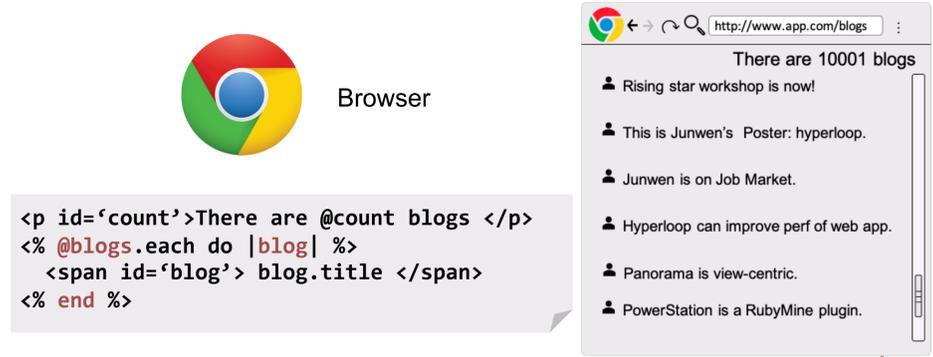
Secret sauce of Hyperloop

- Performance anti-pattern taxonomy based on empirical study.
- Automated detection and optimization of anti-patterns using cross-stack database-aware program analysis



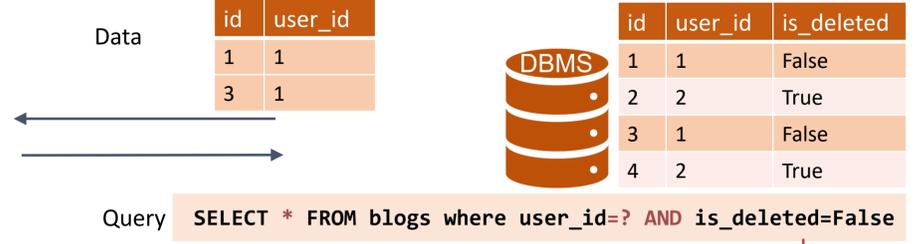
HyperLoop result highlights

- 1000+ performance issues detected and fixed.
- Median 3.1X page speed up, up to 41X.
- Plugin downloaded more than 300 times.



```

def index
  @count = Blog.join(User).count
  @blogs = Blog.where(user_id = ?)
  @blogs.select do |blog|
    !BlockList.all.include?blog.title
  end
end
    
```

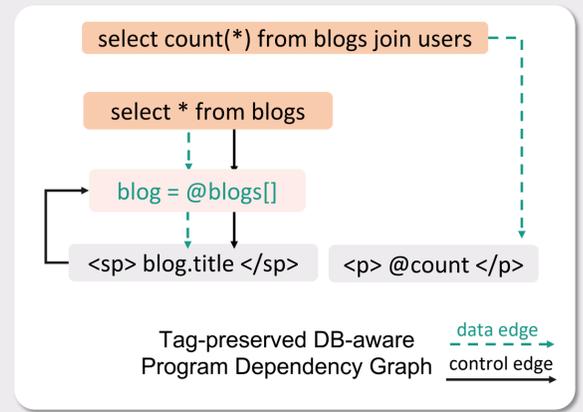


Panorama: View Design Optimization [1,3,5]

17X speed up on page latency!

- ❖ Problem:**
- Developers do not understand the cost behind each HTML tag.
 - Existing optimization does not consider view changes.

- ❖ Solution:**
- A cross-stack cost estimator for every HTML tag.
 - Identify the queries used to generate HTML tag.
 - Estimate query cost.
 - Render the cost in a browser using heatmap.
 - A cross-stack optimizer for HTML design.
 - Synthesize 4 types of view optimization.
 - Paginating
 - Tag Removal
 - Approximation
 - Asynchronous Loading
 - Integrate into a browser-based editor



- ❖ Results:**
- 149 alternate view-design opportunities found!
 - 1.6X page speed up (median), up to 17x!
 - The alternate design is preferred by real-world users!

PowerStation: ORM-API Usage Optimization [1,2]

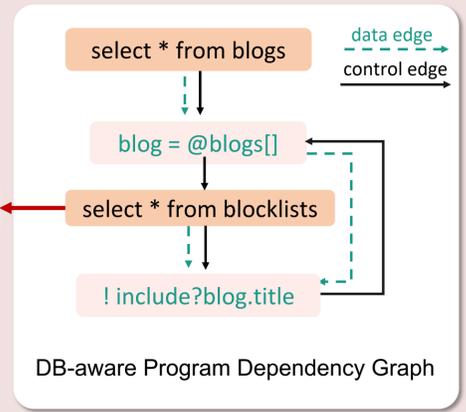
20X speed up on code block!

- ❖ Problem:**
- Developers do not understand the cost difference among ORM APIs.
 - Existing optimization treats ORM APIs as a black box.

- ❖ Solution:**
- A cross-stack program optimizer that understands Database queries behind ORM APIs for:
 - Loop Invariant Motion;
 - Dead Store Elimination;
 - Unused Data-Retrieval Elimination;
 - Common sub-Expression Elimination;
 - Inefficient view rendering elimination
 - Inefficient API replacement;...
 - Integrate into a popular Rails IDE RubyMine as a plugin.

```

blocklists = BlockList.all
@blogs.select do |blog|
  !blocklists.includes?blog.title
end
    
```



- ❖ Results:**
- 1221 inefficiency issues detected, 730 fixed!
 - 1.7X median speed up, up to 20X(code block)!
 - More than 300 download of the plugin!

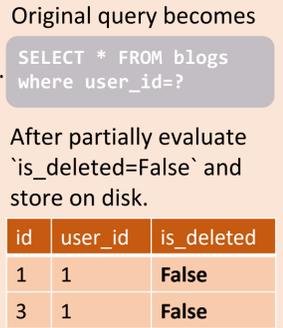
Database Design Optimization [1,4,6]

1.4X speed up on page latency!

- ❖ Problem:**
- Developers do not understand query processing and execution.
 - Database optimizer treats application logic as a black box

- ❖ Solution:**
- A cross-stack program analyzer that identifies:
- Queries benefits from customized physical DB layout.
 - queries with constant parameters in predicate;
 - queries selecting subset of fields;
 - queries joining two or more tables;...
 - Queries benefits from data constraints
 - queries with DISTINCT keyword on unique field;...

- ❖ Results:**
- > 61% of the queries can be optimized!
 - Page speed up 41X(max)!



References:

- Yang, Junwen., Yan, C., Subramaniam, P., Lu, S., & Cheung, A. How not to structure your database-backed web applications: a study of performance bugs in the wild. *ICSE'18. Featured on Hacker News, Morning Paper.*
- Yang, Junwen., Yan, C., Subramaniam, P., Lu, S., & Cheung, A. PowerStation: Automatically detecting and fixing inefficiencies of database-backed web applications in IDE. *FSE'18.*
- Yang, Junwen., Yan, C., Wan, C., Lu, S., & Cheung, A. View-Centric Performance Optimization for Database-Backed Web Applications. *ICSE'19. SIGSOFT Distinguished Paper. Featured on Morning Paper.*
- Yan, C., Cheung, A., Yang, Junwen., & Lu, S. Understanding Database Performance Inefficiencies in Real-world Web Applications. *CIKM'17*
- Yan, C., Cheung, A., Yang, Junwen., & Lu, S. View-Driven Optimization of Database-Backed Web Applications. *CIDR'20*
- Yang, Junwen., Sethi, U., Yan, C., Cheung, A., & Lu, S. Managing data constraints in database-backed web applications. *ICSE'20*